Attorney Docket No. 42P14932

APPLICATION FOR UNITED STATES LETTERS PATENT

For

TECHNIQUE FOR IMPLEMENTING A SECURITY ALGORITHM

Inventor:

Wajdi K. Feghali

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Seventh Floor
Los Angeles, CA 90025-1026
(408) 720-8300

"Express Mail" mailing label number: EV336588048US
Date of Deposit: December 4, 2003
I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231

Christopher P. Marshall
(Typed or printed name of person mailing paper or fee)

C.P. Marshall
(Signature of person mailing paper or fee)

12/4/2003
(Date signed)

TECHNIQUE FOR IMPLEMENTING A SECURITY ALGORITHM

FIELD

[0001] Embodiments of the invention relate to network security algorithms. More particularly, embodiments of the invention relate to the performance of the hash algorithms, including, for example, the secure hash algorithms ("SHA") SHA-1, SHA-128, SHA-192, and SHA-256, as well as message digest (MD) algorithms, such as the MD5 algorithm, within network processor architectures.

BACKGROUND

[0002] Security algorithms may be used to encode or decode data transmitted or received in a computer network through techniques, such as compression.

[0003] In some instances, the network processor may compress or decompress the data in order to help secure the integrity and/or privacy of the information being transmitted or received within the data. The data can be compressed or decompressed by performing a variety of different algorithms, such as hash algorithms.

[0004] One such hash algorithm is the secure hash algorithm 1 ("SHA-1") security algorithm. The SHA-1 algorithm can be a laborious and resource-consuming task for many network processors, however, as it requires numerous mathematically intensive computations within a main recursive compression loop. Moreover, the main compression loop may be performed numerous times in order to compress or decompress a particular amount of data.

[0005] In general, hash algorithms are algorithms that take a large group of data and reduce it to a smaller representation of that data. Hash algorithms may

be used in such applications as security algorithms to protect data from corruption or detection. The SHA-1 algorithm, for example, may reduce groups of 64 bytes of data to 20 bytes of data. Other hash algorithms, such as the SHA-128, SHA-129, and message digest 5 (MD5) algorithms may also be used to reduce large groups of data to smaller ones. Hash algorithms, in general, can be very taxing on computer system performance as the algorithm requires intensive mathematical computations in a recursive main compression loop that is performed iteratively to compress or decompress groups of data.

[0006] Adding to the difficulty in performing the hash algorithms at high frequencies are the latencies, or "bottlenecks," that can occur between operations of the algorithm due to data dependencies between the operations. When performing the algorithm on typical processor architectures, the operations must be performed in substantially sequential fashion because typical processor architectures perform the operations of each iteration of the main compression loop on the same logic units or group of logic units. As a result, if dependencies exist between the iterations of the main loop, a bottleneck forms while unexecuted iterations are delayed to allow the hardware to finish processing the earlier operations.

[0007] These bottlenecks can be somewhat abrogated by taking advantage of instruction-level parallelism ("ILP") of instructions within the algorithm and performing them in parallel execution units.

[0008] Typical prior art parallel execution unit architectures used to perform hash algorithms have had marginal success. This is true, in part, because the

instruction and sub-instruction operations associated with typical hash algorithms rarely have the necessary ILP to allow true independent parallel execution. Furthermore, earlier architectures do not typically schedule operations in such a way as to minimize the critical path associated with long dependency chains among various operations.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] Embodiments of the invention are illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

[0010] Figure 1 illustrates a processor architecture in which one embodiment of the invention may be used.

[0011] Figure 2 illustrates a system in which one embodiment of the invention may be used.

[0012] Figure 3 illustrates an technique for performing a hash algorithm in a pipelined architecture according to one embodiment of the invention.

[0013] Figure 4 illustrates a method for performing a hash algorithm according to one embodiment of the invention.

DETAILED DESCRIPTION

[0014] Embodiments of the invention described herein relate to network security algorithms. More particularly, embodiments of the invention described herein relate to a technique that may be used to improve the performance of the hash algorithms without incurring significant cost.

[0015] At least one embodiment of the invention may be used to improve the performance of hash algorithms by performing various operations associated with the algorithm concurrently, or "pipelining" the operations, within a microprocessor. Pipelining the SHA-1 algorithm, for example involves performing various iterations of the main compression loop at different stages of a microprocessor concurrently in one embodiment of the invention. The extent to which iterative operations of the algorithm may be performed depends, at least in part, on the instruction level parallelism (ILP) of the microprocessor in which the algorithm is executed.

[0016] Figure 1 illustrates a processor architecture in which one embodiment of the invention may be used to perform a hash algorithm while reducing performance degradation, or "bottlenecks," within the processor. In the embodiment of the invention illustrated in Figure 1, the pipeline architecture of the encryption portion 105 of the network processor 100 may operate at frequencies at or near the operating frequency of the network processor itself or, alternatively, at an operating frequency equal to that of one or more logic circuits within the network processor.

[0017] Figure 2 illustrates a computer network in which an embodiment of the invention may be used. The host computer 225 may communicate with a client computer 210 or another host computer 215 by driving or receiving data upon the bus 220. The data is received and transmitted across a network by a program running on a network processor embedded within the network computers. At least one embodiment of the invention 205 may be implemented within the host computer in order to compress that data that is sent to the client computer(s).

[0018] For example, in one embodiment of the invention, a hash algorithm operates on blocks of 512 bits of data at time. In such an embodiment, algorithm compression loop receives a 160-bit state input along with 512 bits of data and produces a 160-bit state output by producing an intermediate state output that is a function of the state input and data input and adds this to the state input in order to produce the 160-bit final state output.

[0019] Various processor architectures having various operating frequencies may be used to facilitate the expedient compression or decompression of the data using a hash algorithm. In one embodiment of the invention, hash algorithms, such as the SHA-1, SHA-128, SHA-192, SHA-256, and the message digest 5 (MD5) algorithms, are performed on pipelined processor architectures operating at frequencies up to, and in excess of, 1.4 GHz.

[0020] For at least one embodiment of the invention, the SHA-1 algorithm can be performed with fewer performance bottlenecks and at greater operating frequencies by taking advantage of the recursive nature of the inner loop of the algorithm, which performs multiple iterations of the equation:

$$TEMP = R^5(A) + F_t(B,C,D) + E + X_t + K_t$$

Where:

$E = D;$

$D = C;$

$C = R^{30}(B);$

$B = A;$

$A = TEMP;$

[0021] A, B, C, D are chaining variables that changes state with each iteration of the loop, such that the function F_t produces a new result for each iteration. The function F_t represents a mathematical operation or operations that performs the SHA-1 algorithm. The function $R^x(v)$ is a left rotation function of x bits by a number of bit positions, v. The rotation function may be implemented using various logic devices, including a shifter. The equation is executed t number of times in order to process X_t units of data Furthermore, the constant, K, may change periodically.

[0022] For example, in one embodiment of the invention, the loop executes 80 times to process 512 bits of data, and the constant, K, changes every 20 iterations of the loop.

[0023] The inner compression loop that executes the above equation may be performed much faster by performing the loop in a pipelined processor architecture, wherein each iteration of the loop is performed by a dedicated pipeline stage or stages. For example, Figure 3 illustrates some of the pipeline

stages and operations that are performed to execute the above equation 80 times in order to process 512 bits of data.

[0024] In pipeline stage 1 305, the constant, K, is added to the first data word 301, X_1, and the chaining variables, B, C, D, and E, are set to their initial state 302. In pipeline stage 2 310, the result of stage 1 is added to the chaining variable E_0 306, chaining variables, B, C, and D, are applied to the function, F 307, and the constant, K, is added to the data word, X_2 308. In pipeline stage 3 320, the result of the function F is added 311 to the sum of E and the result from stage 1, the constant, K, is added to the data word X_3 313, the stage 2 result is added to the chaining variable D 314, and the function is applied to chaining variables A, B, and C 315 after they have been rotated to the left by 30 bits.

[0025] Because the inner loop is executed 80 times in order to process 512 bits, the pipeline architecture illustrated in Figure 3 may require 83 stages to perform the all 80 iterations of the equation. Furthermore, the SHA-1 algorithm may require that the initial states of the chaining variables be added to the final states of the chaining variables, which requires an extra 5 pipeline stages, for a total of 88 pipeline stages to completely process the 512 bits of data.

[0026] In other embodiments, fewer pipeline stages may be used to process the 512 bits by performing several operations at each pipeline stage. Furthermore, fewer or more pipeline stages may be used in other embodiments of the invention depending upon the hash algorithm to be performed. The determination of what pipeline stages perform which operations associated with each iteration of the equation is largely determined by scheduling logic that

attempts to schedule these operations according to their data dependencies. In this manner, the scheduler can use the pipeline stages in the most efficient manner possible, thereby preventing lengthy bottlenecks.

[0027] For example, if an operation performed in iteration 2 of the inner loop is dependent upon data from an operation performed in iteration 1, the second pipeline may remain partially idle for a time until the data from iteration 1 (performed in the first pipeline stage) is available. However, if these operations are performed in parallel with operations having similar data dependencies and therefore similar delay, the scheduler can perform these operations at one time at different pipeline stages. As a result, bottlenecks and delays incurred by data dependencies are minimized, allowing the pipeline architecture to operate at higher frequencies limited only by the operating frequency of the processor architecture or hardware therein, such as the adder circuits.

[0028] Figure 4 illustrates a method for carrying out the invention according to one embodiment. The main loop equation of the SHA-1 must be decoded into separate operations corresponding to the operations that must be performed at each iteration of the loop at operation 405. Operations corresponding to each iteration of the loop may then be scheduled for execution within a particular pipeline stage at operation 410.

[0029] The choice of which stage in which to perform the operation(s) depends, at least in part, on the dependencies between the operations. In order to perform the algorithm at the highest performance level possible, the critical paths (operations having lengthy dependency chains) must be found and

scheduled in such a manner so as to impose a minimum amount of delay on the performance of other operations and the algorithm in general.

[0030] Many of the scheduled operations may then be performed by the various pipeline stages in parallel at operation 415 if they have few or no dependencies from earlier operations.

[0031] In addition to performing the operations corresponding to each iteration of the main compression loop, the initial and final states of the chaining variables must be added to each other in order to produce the final output state at operation 420.

[0032] Embodiments of the invention may be performed using logic consisting of standard complementary metal-oxide-semiconductor ("CMOS") devices (hardware) or by using instructions (software) stored upon a machine-readable medium, which when executed by a machine, such as a processor, cause the machine to perform a method to carry out the steps of an embodiment of the invention. Alternatively, a combination of hardware and software may be used to carry out embodiments of the invention.

[0033] While the invention has been described with reference to illustrative embodiments, this description is not intended to be construed in a limiting sense. Various modifications of the illustrative embodiments, as well as other embodiments, which are apparent to persons skilled in the art to which the invention pertains are deemed to lie within the spirit and scope of the invention.